arXiv:1604.04894v1 [cs.AI] 17 Apr 2016

A global constraint for closed itemset mining

M. Maamar[a,b], N. Lazaar[b], S. Loudni[c], and Y. Lebbah[a]

(a) University of Oran 1, LITIO, Algeria.
(b) University of Montpellier, LIRMM, France.
(c) University of Caen, GREYC, France.

Abstract. Discovering the set of closed frequent patterns is one of the fundamental problems in Data Mining. Recent Constraint Programming (CP) approaches for declarative itemset mining have proven their usefulness and flexibility. But the wide use of reified constraints in current CP approaches raises many difficulties to cope with high dimensional datasets. This paper proposes CLOSEDPATTERN global constraint which does not require any reified constraints nor any extra variables to encode efficiently the Closed Frequent Pattern Mining (CFPM) constraint. CLOSEDPATTERN captures the particular semantics of the CFPM problem in order to ensure a polynomial pruning algorithm ensuring domain consistency. The computational properties of our constraint are analyzed and their practical effectiveness is experimentally evaluated.

1 Introduction

Frequent Pattern Mining is a well-known and the most popular research field of data mining. Originally introduced by [1], it plays a key role in many data mining applications. These applications include the discovery of frequent itemsets and association rules [1], correlations [3] and many other data mining tasks.

In practice, the number of frequent patterns produced is often huge and can easily surpass the size of the input dataset. Based on this statement, it was important to identify a condensed representation of frequent patterns. On the other hand, most of frequent patterns are redundant where it is possible to derive them from other found patterns. That is, *closed frequent patterns* are one of the concise and condensed representations avoiding redundancy.

Discovering the set of closed frequent patterns is one of the fundamental problems in Data Mining. Several specialized approaches have been proposed to discover closed frequent patterns (e.g., A-CLOSE algorithm [13], `CHARM` [18], `CLOSET` [14], `LCM` [16]).

Over the last decade, the use of Constraint Programming paradigm (CP) to model and to solve Data Mining problems has received a considerable attention [5,7,10]. The declarative aspect represents the key success of the proposed CP approaches. Doing so, one can add/remove any user-constraint without the need of developing specialized resolution methods.

Relating to the Closed Frequent Pattern Mining problem (CFPM), Luc De Raedt et. al., propose to express the different constraints that we can have in Pattern Mining as a CP model [7]. The model is expressed on boolean variables representing items and transactions, with a set of reified sums as constraints. The drawback is the wide use of reified constraints in the CP model, which makes the scalability of the approach questionable.

This paper proposes CLOSEDPATTERN global constraint which does not require any reified constraints nor any extra variables to encode efficiently the Closed Frequent Pattern Mining (CFPM) constraint. CLOSEDPATTERN captures the particular semantics of the CFPM problem in order to ensure a polynomial pruning algorithm ensuring domain consistency.

Experiments on several known large datasets show that our approach clearly outperforms CP4IM [5] and achieves scalability while it is a major issue for CP approaches. These experiments also show that the fewer the number of closed patterns, the better is the performance of CLOSEDPATTERN. This is an expected result of the fact that CLOSEDPATTERN insures domain consistency.

The paper is organized as follows. Section 2 recalls preliminaries. Section 3 provides a critical review of specialized methods and CP approaches for CFPM. Section 4 presents the global constraint CLOSEDPATTERN. Section 5 reports experiments we performed. Finally, we conclude and draw some perspectives.

2 Background

In this section, we introduce some useful notions in closed frequent pattern mining and constraint programming.

2.1 Closed frequent pattern mining

Let $\mathcal{I} = \{1, ..., n\}$ be a set of n *items* identifiers and $\mathcal{T} = \{1, ..., m\}$ a set of *transactions* identifiers. A pattern p (i.e., itemset) is a subset of $\mathcal{I}$. The language of patterns corresponds to $\mathcal{L}_{\mathcal{I}} = 2^{\mathcal{I}}$. A transaction database is a set $\mathcal{D} \subseteq \mathcal{I} \times \mathcal{T}$. The set of items corresponding to a transaction identified by t is denoted $\mathcal{D}[t] = \{i \mid (i, t) \in \mathcal{D}\}$. A transaction t is an occurrence of some pattern p iff the set $\mathcal{D}[t]$ contains p (i.e. $p \subseteq \mathcal{D}[t]$).

The cover of p, denoted by $\mathcal{T}_{\mathcal{D}}(p)$, is the set of transactions containing p (i.e. $\mathcal{T}_{\mathcal{D}}(p) = \{t \in \mathcal{T} \mid p \subseteq \mathcal{D}[t]\}$). Given $S \subseteq \mathcal{T}$ a subset of transactions, $\mathcal{I}_{\mathcal{D}}(S) = \bigcap_{t \in S} \mathcal{D}[t]$ is the set of common items of S. The (absolute) frequency of a pattern p is the size of its cover (i.e., $freq_{\mathcal{D}}(p) = |\mathcal{T}_{\mathcal{D}}(p)|$). Let $\theta \in \mathbb{N}^+$ be some given constant called a *minimum support*. A pattern p is frequent if $freq_{\mathcal{D}}(p) \geq \theta$.

Example 1. Consider the transaction database in Table 1a. We have $\mathcal{T}_{\mathcal{D}}(CE) = \{2, 3, 5, 6\}$, $freq_{\mathcal{D}}(CE) = 4$ and $\mathcal{I}_{\mathcal{D}}(\{2, 3, 5, 6\}) = BCE$.

The closure of a pattern p in $\mathcal{D}$ is the set of common items of its cover $\mathcal{T}_{\mathcal{D}}(p)$, which is denoted $Clos(p) = \mathcal{I}_{\mathcal{D}}(\mathcal{T}_{\mathcal{D}}(p))$. A pattern is closed $closed_{\mathcal{D}}(p)$ if and only if $Clos(p) = p$.

Definition 1 (Closed Frequent Pattern Mining (CFPM)). *Given a transaction database $\mathcal{D}$ and a minimum support threshold θ. The closed frequent pattern mining problem is the problem of finding all patterns p such that $(freq_{\mathcal{D}}(p) \geq \theta)$ and $(Clos(p) = p)$.*

Example 2. For $\theta = 2$, the set of closed frequent patterns in Table 1a is $C\langle 5\rangle$[1], $BE\langle 5\rangle$, $BCE\langle 4\rangle$, $ABCE\langle 2\rangle$ and $AC\langle 2\rangle$.

[1] Value between $\langle . \rangle$ indicates the frequency of a pattern.

t	Items
t_1	A C D
t_2	B C E
t_3	A B C E
t_4	B E
t_5	A B C E
t_6	B C E

(a)

t	A	B	C	D	E
t_1	1	0	1	1	0
t_2	0	1	1	0	1
t_3	1	1	1	0	1
t_4	0	1	0	0	1
t_5	1	1	1	0	1
t_6	0	1	1	0	1

(b)

Table 1: A transaction database $\mathcal{D}$ (a) and its binary matrix (b).

Closed frequent patterns provide a minimal representation of frequent patterns , i.e., we can derive all frequent patterns with their exact frequency value from the closed ones.

Search Space Issues. In pattern mining, the search space contains $2^{\mathcal{I}}$ candidates. Given a large number of items $\mathcal{I}$, a naïve search that consists of enumerating and testing the frequency of pattern candidates in a given dataset is infeasible. The main property exploited by most algorithms to reduce the search space is that frequency is *monotone decreasing* with respect to extension of a set.

Property 1 (Anti-monotonicity of the frequency). Given a transaction database $\mathcal{D}$ over $\mathcal{I}$, and two patterns $X, Y \subseteq \mathcal{I}$. Then, $X \subseteq Y \rightarrow freq_{\mathcal{D}}(Y) \leq freq_{\mathcal{D}}(X)$.

Hence, any subset (resp. superset) of a frequent (resp. infrequent) pattern is also a frequent (resp. infrequent) pattern.

2.2 CFPM under constraints

Constraint-based pattern mining aims at extracting all patterns p of $\mathcal{L}_{\mathcal{I}}$ satisfying a query $q(p)$ (conjunction of constraints), which usually defines what we call a *theory* [11]: $Th(q) = \{p \in \mathcal{L}_{\mathcal{I}} \mid q(p)\ is\ true\}$. A common example is the frequency measure leading to the minimal frequency constraint. It is also possible to have other kind of (user-)constraints. For instance, Constraints on the size of the returned patterns, $minSize(p, \ell_{min})$ constraint (resp. $maxSize(p, \ell_{min})$) holds if the number of items of p is greater or equal (resp. less or equal) to ℓ_{min}. Constraints on the presence of an item in a pattern $item(p, i)$ that states that an item i must be (or not) in a pattern p.

2.3 Dataset representations

The algorithms for frequent pattern mining differ mainly on the way that the dataset is represented.

Horizontal representation $\mathcal{H}$. Here, the transaction dataset is represented as a list of transaction. The Apriori algorithm [1] is one of the approaches that adopt this obvious representation. The drawback of such representation is the fact that we need several passes to update the support of pattern candidates.

Vertical representation $\mathcal{V}$. This representation uses a list of items where for each item, we have the list of transactions where it appears. Many algorithms adopt this

representation [19,17,4]. The key advantage of using such representation is that the support of a pattern candidate can easily be obtained by intersecting the lists of its items.

Hybrid representation $\mathcal{HV}$. The transaction dataset here is represented dually, horizontally and vertically. It is successfully used within LCM [16] and FPgrowth [9] algorithms.

2.4 CSP and Global Constraints

A *Constraint Satisfaction Problem* (CSP) consists of a set X of n variables, a domain $\mathcal{D}$ mapping each variable $X_i \in X$ to a finite set of values $D(X_i)$, and a set of constraints $\mathcal{C}$. An assignment σ is a mapping from variables in X to values in their domains. A constraint $c \in \mathcal{C}$ is a subset of the cartesian product of the domains of the variables that are in c. The goal is to find an assignment such that all constraints are satisfied.

Domain consistency (DC). Constraint solvers typically use backtracking search to explore the search space of partial assignments. At each assignment, filtering algorithms prune the search space by enforcing local consistency properties like domain consistency. A constraint c on X is domain consistent, if and only if, for every $X_i \in X$ and every $d_i \in D(X_i)$, there is an assignment σ satisfying c such that $X_i = d_i$.

Global constraints are constraints capturing a relation between a non-fixed number of variables. These constraints provide the solver with a better view of the structure of the problem. Examples of global constraints are AllDifferent, Regular and Among (see [15]). Global constraints cannot be efficiently propagated by generic local consistency algorithms, which are exponential in the number of the variables of the constraint. Dedicated filtering algorithms are constructed to achieve polynomial time complexity in the size of the input, i.e., the domains and extra parameters. This is the aim of this paper, which proposes a filtering algorithm for the frequent closed pattern constraint.

3 Related works

This section provides a critical review of ad specialized methods and CP approaches for CFPM.

Specialized methods for CFPM. `CLOSE` [13] was the first algorithm proposed to extract closed frequent patterns (CFPs). It uses an Apriori-like bottom-up method. Later, Zaki and Hsiao [18] proposed a depth-first algorithm based on a vertical database format e.g. `CHARM`. In [14], Pei et al. extended the FP-growth method to a method called `CLOSET` for mining CFPs. Finally, Uno et al. [16] have proposed `LCM`, one of the most fastest frequent itemset mining algorithm. It employs a hybrid representation based on vertical and horizontal ones.

CP methode for itemset mining. Luc De Raedt et al. have proposed in [5] a CP model for itemset mining (`CP4IM`). They show how some constraints (e.g. frequency, maximality, closedness) can be modeled as CSP [12,8]. This modeling uses two sets of boolean variables M and T: (1) item variables $\{M_1, M_2, ..., M_n\}$ where, given a pattern P, $(M_i = 1)$ iff $(i \in P)$; (2) transaction variables $\{T_1, T_2, ..., T_m\}$ where $(T_t = 1)$ iff $(P \subseteq t)$.

The relationship between M and T is modeled by reified constraints stating that, for each transaction t, $(T_t = 1)$ iff M is a subset of t. A great consequence is that the encoding of the frequency measure is straightforward: $freq_{\mathcal{D}}(M) = \sum_{t \in T} T_t$. But such an encoding has a major drawback since it requires $(m = \#\mathcal{T})$ reified constraints to encode the whole database. This constitutes a strong limitation of the size of the databases that could be managed.

We propose in the next section the CLOSEDPATTERN global constraint to encode both the minimum frequency constraint and the closedness constraint. This global constraint does not require any reified constraints nor any extra variables.

4 CLOSEDPATTERN global constraint for CFPM

This section presents the CLOSEDPATTERN global constraint for the CFPM problem.

4.1 Consistency checking and filtering

Let P be the unknown pattern we are looking for. The unknown pattern P is encoded with boolean item variables $\{P_1, ..., P_n\}$, where $D(P_i) = \{0, 1\}$ and $(P_i = 1)$ iff $i \in P$. Let σ be a partial assignment of variables P. σ can be partitioned into three distinct subsets:

- present items $\sigma^+ = \{i \in \mathcal{I} \,|\, D(P_i) = \{1\}\}$,
- absent items $\sigma^- = \{i \in \mathcal{I} \,|\, D(P_i) = \{0\}\}$,
- free items $\sigma^* = \{i \in \mathcal{I} \,|\, D(P_i) = \{0, 1\}\}$.

The global constraint CLOSEDPATTERN ensures both minimum frequency constraint and closedness constraint.

Definition 2 (CLOSEDPATTERN global constraint). *The* CLOSEDPATTERN$(P, \mathcal{D}, \theta)$ *constraint holds if and only if there exists an assignment* $\sigma = \langle d_1, ..., d_n \rangle$ *of variables* P *such that* $freq_{\mathcal{D}}(\sigma^+) \geq \theta$ *and* $closed_{\mathcal{D}}(\sigma^+)$.

Example 3. Consider the transaction database of Table 1a with $\theta = 2$. Let $P = \langle P_1, \ldots, P_5 \rangle$ with $D(P_i) = \{0, 1\}$ for $i = 1..5$. Consider the closed pattern BCE encoded by $P = \langle 01101 \rangle$, where $\sigma^+ = \{B, C, E\}$ and $\sigma^- = \{A, D\}$. CLOSEDPATTERN$(P, \mathcal{D}, 2)$ holds since $freq_{\mathcal{D}}(\{B, C, E\}) \geq 2$ and $closed_{\mathcal{D}}(\{B, C, E\})$.

Definition 3 (Full extension item [2]). *Let* $\sigma = \langle d_{i_1}, \ldots, d_{i_\ell} \rangle$ *be a partial assignment of* ℓ *variables* $\langle P_{i_1}, \ldots, P_{i_\ell} \rangle$*, and* j *an item such that* $j \notin \sigma^+$*. The item* j *is called a full extension of* σ *iff* $\mathcal{T}_{\mathcal{D}}(\sigma^+) = \mathcal{T}_{\mathcal{D}}(\sigma^+ \cup \{j\})$.

Let σ be a partial assignment of variables P and i a free item. We denote by $\mathcal{V}_{\mathcal{D}}^{\sigma^+}(i)$ the cover of item i within the current cover of a pattern σ^+:

$$\mathcal{V}_{\mathcal{D}}^{\sigma^+}(i) = \mathcal{T}_{\mathcal{D}}(\sigma^+ \cup \{i\}) = \mathcal{T}_{\mathcal{D}}(\sigma^+) \cap \mathcal{T}_{\mathcal{D}}(\{i\}).$$

We first show when a partial assignment is consistent with respect to CLOSEDPATTERN constraint.

Proposition 1. *Let $\sigma = \langle d_{i_1}, \ldots, d_{i_l} \rangle$ be a partial assignment of ℓ variables $\langle P_{i_1}, \ldots, P_{i_\ell} \rangle$. We say that σ is a consistent partial assignment iff: $freq_{\mathcal{D}}(\sigma^+) \geq \theta$ and $\nexists j \in \sigma^-$ s.t. j is a full extension of σ.*

Proof:

According to the anti-monotonicity property of the frequency (cf. property 1), if the partial assignment σ is infrequent (i.e., $freq_{\mathcal{D}}(\sigma^+) < \theta$), it cannot, under any circumstances, be extended to a closed pattern.

Given now a frequent partial assignment σ (i.e., $freq_{\mathcal{D}}(\sigma^+) \geq \theta$), let us take $j \in \sigma^-$ s.t. j is a full extension of σ. It follows that $\mathcal{T}_{\mathcal{D}}(\sigma^+) = \mathcal{T}_{\mathcal{D}}(\sigma^+ \cup \{j\}) = \mathcal{V}_{\mathcal{D}}^{\sigma^+}(j)$. Therefore, $Clos(\sigma^+) = Clos(\sigma^+ \cup \{j\})$. Since σ^+ without j (j being in σ^-) cannot be extended to a closed pattern, the result follows. If there is no item $j \in \sigma^-$ s.t. j is a full extension of σ, then the current assignment σ can be definitely extended to a closed itemset by adopting all the full extension items to form a closed pattern. □

We now give the CLOSEDPATTERN filtering rules by showing when a value of a given variable is inconsistent.

Proposition 2 (CLOSEDPATTERN Filtering rules). *Let $\sigma = \langle d_{i_1}, \ldots, d_{i_\ell} \rangle$ be a consistent partial assignment of ℓ variables $\langle P_{i_1}, \ldots, P_{i_\ell} \rangle$, and P_j ($j \in \sigma^*$) be a free variable. The following two cases characterize the inconsistency of the values 0 and 1 of P_j:*

- $0 \notin D(P_j)$ *iff: j is a full extension of σ.* *(rule 1)*
- $1 \notin D(P_j)$ *iff:* $\begin{cases} |\mathcal{V}_{\mathcal{D}}^{\sigma^+}(j)| < \theta \quad \vee & \textit{(rule 2)} \\ \exists k \in \sigma^-, \mathcal{V}_{\mathcal{D}}^{\sigma^+}(j) \subseteq \mathcal{V}_{\mathcal{D}}^{\sigma^+}(k). & \textit{(rule 3)} \end{cases}$

Proof:

Let σ be a consistent partial assignment and P_j be a free variable.

$0 \notin D(P_j)$: ($\Rightarrow$) Let 0 be an inconsistent value. In this case, P_j can only take value 1. It means that $Clos(\sigma^+) = Clos(\sigma^+ \cup \{j\})$. Thus, $\mathcal{T}_{\mathcal{D}}(\sigma^+) = \mathcal{T}_{\mathcal{D}}(\sigma^+ \cup \{j\})$. By definition 3, j is a full extension of σ.
($\Leftarrow$) Let j be a full extension of σ, which means that $Clos(\sigma^+) = Clos(\sigma^+ \cup \{j\})$ (def. 3). The value 0 is inconsistent where j cannot be in σ^- (proposition 1).

$1 \notin D(P_j)$: ($\Rightarrow$) Let 1 be an inconsistent value. This can be the case if the frequency of the current pattern σ^+ is set up bellow the threshold θ by adding the item j (i.e., $|\mathcal{V}_{\mathcal{D}}^{\sigma^+}(j)| < \theta$). Or, $\sigma^+ \cup \{j\}$ cannot be extended to a closed itemset: this is the case when it exists an item k such that at each time the item j belongs to a transaction in the database, k belongs as well ($\mathcal{V}_{\mathcal{D}}^{\sigma^+}(j) \subseteq \mathcal{V}_{\mathcal{D}}^{\sigma^+}(k)$). Conversly, the lack of k (i.e., $k \in \sigma^-$) implies the lack of j as well. This means that: ($P_k = 0 \Rightarrow P_j = 0$).
($\Leftarrow$) This is a direct consequence of proposition 1. □

Example 4. In line of example 3, consider a partial assignment σ s.t. $\sigma^+ = \{B\}$ and $\sigma^- = \emptyset$. CLOSEDPATTERN$(P, \mathcal{D}, \theta)$ will remove value 1 from $D(P_4)$ (item D) and value 0 from $D(P_5)$ (item E) since, resp., $|\mathcal{V}_{\mathcal{D}}^{\sigma^+}(D)| < 2$ and E is a full extension of

σ. Now, we have $\sigma^+ = \{B, E\}$ and $\sigma^- = \{D\}$. Suppose that the variable P_3 is set to 0 (item C). Again, CLOSEDPATTERN$(P, \mathcal{D}, \theta)$ will remove value 1 from $D(P_1)$ (item A) since the lack of C implies the lack of A in $\mathcal{D}$ (i.e., $\mathcal{V}_{\mathcal{D}}^{\sigma^+}(A) \subseteq \mathcal{V}_{\mathcal{D}}^{\sigma^+}(C)$).

4.2 CLOSEDPATTERN Filtering Algorithm

In this section, we present an algorithm enforcing domain consistency for CLOSEDPATTERN constraint. Algorithm 1 maintains the consistency based on the specificity of the CFPM problem (see proposition 1 and 2). CLOSEDPATTERN is considered as a global constraint since all variables share the same internal data structures that awake and drive the filtering.

FILTER-CLOSEDPATTERN algorithm exploits internal data structures enabling to enhance the filtering process. At each call, we maintain incrementaly $\sigma =< \sigma^+, \sigma^-, \sigma^* >$ and the cover of σ^+ (i.e., $\mathcal{T}_{\mathcal{D}}(\sigma^+)$). Using these two structures, one can check if an item is present or not in the vertical dataset $\mathcal{V}_{\mathcal{D}}$.

Algo.1 takes as input the vertical dataset $\mathcal{V}_{\mathcal{D}}$, a minimum support threshold θ, the item k of the last assigned variable P_k, the current partial assignment σ where $\sigma^* \neq \emptyset$, and the variables P. As output, algo.1 will reduce de domain of P_i and therefore, increase σ^+ and/or σ^-, and decrease σ^*. The algorithm starts by checking if the last variable P_k is instantiated to 0 (line 3). In such case, the item k is added to (resp. removed from) σ^- (resp. σ^*) (line 4). Afterwards, we check if the lack of the item k induces the lack of other free items of σ^* (lines 6-8), which corresponds to rule 3 of proposition 2. Otherwise, P_k is assigned to 1 or we are in the level of the first call of the FILTER-CLOSEDPATTERN where no variable is instantiated. Here, we add to (resp. remove from) σ^+ (resp. σ^*) the item k if $P_k = 1$ (line 10). Now, the first thing to check is to verify if the current assignement is consistent or not (proposition 1). The line 10 is doing so by checking if the frequency of the current assignment (i.e., the size of the cover $\mathcal{T}_{\mathcal{D}}(\sigma^+)$) is below the threshold θ.

The lines from 11 to 21 represents the application of proposition 2 on the remaining free items σ^*. First, lines 12 to 14 prune the value 0 from each item that is a full extension of the current assignment (rule 1 of prop.2). Second, lines 15 to 17 prune the value 1 from the infrequent items (rule 2 of prop.2). Finaly, lines 18 to 21 implement the rule 3 of prop.2 and prune the value 1 from each free item $i \in \sigma^*$ s.t. its cover is a superset of the cover of an absent item $j \in \sigma^-$ ($\mathcal{V}_{\mathcal{D}}^{\sigma^+}(i) \subseteq \mathcal{V}_{\mathcal{D}}^{\sigma^+}(j)$).

Theorem 1. *Given a transaction database $\mathcal{D}$ of n items and m transactions, and a threshold minsup θ. Algorithm* FILTER-CLOSEDPATTERN *establishes domain consistency on the* CLOSEDPATTERN *constraint, or proves that it is inconsitent in $O(n^2 \times m)$ with a space complexity of $O(n \times m)$.*

Proof:

Since that CLOSEDPATTERN implements exactly the rules given in Proposition 2, it prunes all inconsistent values, and consequently ensures domain consistency (see the given description of Algorithm 1).

Algorithm 1: FILTER-CLOSEDPATTERN ($\mathcal{V}_\mathcal{D}, \theta, k, \sigma, P$)

```
1  Input: V_D : vertical database; θ : minimum support; k : instantiated item
2  InOut: P = {P_1 ... P_n}: boolean item variables; σ : current assignment.
3  begin
4  if P_k = 0 then
5      σ⁻ ← σ⁻ ∪ {k};    σ* ← σ* \ {k}
6      foreach i ∈ σ* : V_D^{σ+}(i) ⊆ V_D^{σ+}(k) do
7          D(P_i) ← D(P_i) − {1};
8          σ⁻ ← σ⁻ ∪ {i};    σ* ← σ* \ {i}
9  else
10     if P_k = 1 then σ+ ← σ+ ∪ {k};    σ* ← σ* \ {k}  if ( |T_D(σ+)| < θ) then
       return false
11     foreach i ∈ σ* do
12         if |V_D^{σ+}(i)| = |T_D(σ+)| then
13             D(P_i) ← D(P_i) − {0};
14             σ+ ← σ+ ∪ {i};    σ* ← σ* \ {i}
15         else if |V_D^{σ+}(i)| < θ then
16             D(P_i) ← D(P_i) − {1}
17             σ⁻ ← σ⁻ ∪ {i};    σ* ← σ* \ {i}
18         else
19             foreach j ∈ σ⁻ : V_D^{σ+}(i) ⊆ V_D^{σ+}(j) do
20                 D(P_i) ← D(P_i) − {1};
21                 σ⁻ ← σ⁻ ∪ {i};    σ* ← σ* \ {i}
22 return true;
```

Let $n = |\mathcal{I}|$ and $m = |\mathcal{T}|$. First, we need to compute $\mathcal{T}_\mathcal{D}(\sigma^+)$ which requires at most $O(n \times m)$. This is done only once. The cover $\mathcal{V}_\mathcal{D}^{\sigma^+}(i)$ can be computed by intersecting $\mathcal{T}_\mathcal{D}(\sigma^+)$ (already computed) and $\mathcal{T}_\mathcal{D}(\{i\})$ (given by the vertical representation) within at most $O(m)$. Checking rule 1 and 2 on all free variables can be done in $O(n \times m)$ (lines 12-17). However, checking rule 3 is quadratic at lines 6-8 (i.e., $O(n \times m)$) and cubic at lines 18-21 (i.e., $O(n \times (n \times m))$), where checking if a cover $\mathcal{V}_\mathcal{D}^{\sigma^+}(i)$ is a subset of another cover can be done in $O(m)$. Finally, the worst case complexity is $O(n \times (n \times m))$.

The space complexity of the filtering algorithm lies in the storage of $\mathcal{V}_\mathcal{D}$, σ and the cover $\mathcal{T}$ data structures. The vertical representation $\mathcal{V}_\mathcal{D}$ requires at most $n \times m$. In the worst case, we have to store n items within σ and m transactions within $\mathcal{T}$. That is, the worst case space complexity is $O(n \times m + n + m) = O(n \times m)$. □

During the solving process in depth first search, the whole space complexity is about $O(n \times (m+n))$ since that: (1) the depth is at most n; (2) σ and $\mathcal{T}$ require $O(n \times (m+n))$; (3) the vertical representation is the same data used all along the solving process $O(n \times m)$; (4) Finally we have $O(n \times (m+n)) + O(n \times m) = O(n \times (m+n))$.

Dataset	$\|\mathcal{T}\|$	$\|\mathcal{I}\|$	$\|\widehat{\mathcal{T}}\|$	ρ	type of data
Chess	3 196	75	37	49%	game steps
Connect	67 557	129	43	33%	game steps
Mushroom	8 124	119	23	19%	species of mushrooms
Pumsb	49 046	7 117	74	1%	census data
BMS-Web-View1	59 601	497	2.5	0.5%	web click stream
T10I4D100K	100 000	1 000	10	1%	synthetic dataset
T40I10D100K	100 000	1 000	40	4%	synthetic dataset
Retail	88 162	16 470	10	0.06%	retail market basket data

Table 2: Dataset Characteristics.

5 Experiments

We made some experiments to evaluate and compare our global constraint CLOSED-PATTERN with the state of the art methods (CP and specialized methods for CFPM). We first present the benchmark datasets we used for our experiments and we give a brief description on the followed protocol.

Benchmark datasets. The reported results are on several real and synthetic datasets [18,6] from FIMI repository[2] with large size. These datasets have varied characteristics and representing different application domains as shown in Table 2.

Table 2 reports for each dataset, the number of transaction $|\mathcal{T}|$, the number of items $|\mathcal{I}|$, the average size of transaction $|\widehat{\mathcal{T}}|$ and its density ρ (i.e., $|\widehat{\mathcal{T}}|/|\mathcal{I}|$). As we can see, we select datasets by varying the number of transactions, the number of items, but also the density ρ. Here we have datasets that are very dense like chess and connect (resp. 49% and 33%), others that are very sparse like Retail and BMS-Web-View1 (resp. 0.06% and 0.5%).

Experimental protocol. The implementation of our approach was carried out in the `Gecode` solver[3]. All experiments were conducted on an Intel Core i5-2400 @ 3.10 GHz with 8 Gb of RAM with a timeout of $3600s$. For each dataset, we varied the $minsup$ threshold until the methods are not able to complete the extraction of all closed patterns within the timeout limit. We compare our approach (CLOSEDPATTERN) with:

1. `CP4IM`, the most popular CP approach for CFPM.
2. `LCM`, the most popular specialized method for CFPM.

We experiment using the available distributions of `LCM-v3` [4] and `CP4IM` [5]. We denote the fact that the underlying solver used by `CP4IM` is `Gecode` solver.

Discussion. First we compare CLOSEDPATTERN with the most efficient CP method `CP4IM`. CPU times of the two methods are given in Fig. 1.

Let us take the mushroom dataset, here CLOSEDPATTERN clearly outperforms `CP4IM`. Moreover, the gains in terms of CPU time becomes significant when the *minsup* threshold decreases. For instance, with a $minsup = 0.05\%$, the CLOSEDPATTERN constraint is about 5 times faster than `CP4IM`. If we take connect and chess datasets, `CP4IM`

[2] `http://fimi.ua.ac.be/data/`
[3] `http://www.gecode.org`
[4] `http://research.nii.ac.jp/~uno/codes.htm`
[5] `https://dtai.cs.kuleuven.be/CP4IM/`

performs better in terms of CPU time. The behavior of CLOSEDPATTERN on the two datasets can be explained by their important density. With such datasets, the number of closed patterns is quite huge, which reduce the pruning power of the global constraint. It is important to stress that we are testing the pruning power of our global constraint without any modification on the search part. On the remaining datasets, `CP4IM` reaches an out of memory state due to huge number of reified constraints. For instance, if we take T40I10D100K dataset, the CP model produced by `CP4IM` contains $|\mathcal{T}| = 100\,000$ reified constraints to express the coverage constraint, $2 \times |\mathcal{I}| = 2 \times 1\,000$ reified constraints to express the closure and frequency constraints. That is, the CP solver has to load in memory a CP model of $102\,000$ reified constraint.

A particular observation that we can make is the fact that our global constraint is able to handle all datasets handled by the specialized algorithm `LCM-v3`, which is not the case of `CP4IM` due to the size of the CP model to load. The result is also quite competitive comparing to `LCM-v3` in terms of CPU time. For instance, if we take the BMS-Web-View1 dataset with $minsup$ = 0.06%, the extraction of all closed patterns is achieved in less than $10s$ with CLOSEDPATTERN for $0.1s$ with `LCM-v3`. The same observation on T10I4D100K dataset (for $minsup \leq 0.5\%$).

To complement the results, The table given in Figure 1 reports a result comparison between CLOSEDPATTERN and `CP4IM`. For each dataset and for each value of $minsup$, we report the number of closed patterns (Col. 3), the number of calls to the propagate routine of `Gecode` (Col. 4), the number of explored nodes of the search tree (Col. 5), and the memory usage (Col. 6).

If we comeback to mushroom, CLOSEDPATTERN explores less nodes than `CP4IM` and thus, we have less calls of the propagator, while on chess and connect we have the opposite. This can be explained by the fact that we do not use any particular heuristic in the search part. The second observation is that our approach is very efficient in terms of number of propagations. For CLOSEDPATTERN, the number of propagations remains meaningless comparing to the one using `CP4IM` on small minsup values. This is due to the huge number of reified constraints used in `CP4IM`. As a last observation relating to memory consumption, CLOSEDPATTERN uses very low amount of memory on mushroom and connect, but consumes a little bit more memory on chess comparing to `CP4IM`.

At the end, comparing CLOSEDPATTERN with `LCM-v3` shows that our global constraint remains competitive knowing that it just enforces domain consistency at each node without any dedicated search heuristic.

6 Conclusion

In this paper, we have proposed the global constraint CLOSEDPATTERN for Closed Frequent Pattern Mining. CLOSEDPATTERN captures the particular semantics of the CFPM problem in order to ensure a polynomial pruning algorithm ensuring domain consistency. Experiments on several known large datasets show that our global constraint is clearly efficient and achieves scalability while it is a major issue for CP approaches.

References

1. Agrawal, R., Imieliński, T., Swami, A.: Mining association rules between sets of items in large databases. In: ACM SIGMOD Record. vol. 22, pp. 207–216. ACM (1993)
2. Borgelt, C.: Frequent item set mining. Wiley Interdisc. Rew.: Data Mining and Knowledge Discovery 2(6), 437–456 (2012), http://dx.doi.org/10.1002/widm.1074
3. Brin, S., Motwani, R., Silverstein, C.: Beyond market baskets: Generalizing association rules to correlations. In: SIGMOD. pp. 265–276 (1997), http://doi.acm.org/10.1145/253260.253327
4. Burdick, D., Calimlim, M., Flannick, J., Gehrke, J., Yiu, T.: MAFIA: A maximal frequent itemset algorithm. IEEE Trans. Knowl. Data Eng. 17(11), 1490–1504 (2005), http://dx.doi.org/10.1109/TKDE.2005.183
5. De Raedt, L., Guns, T., Nijssen, S.: Constraint programming for itemset mining. In: Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining. pp. 204–212. ACM (2008)
6. Grahne, G., Zhu, J.: Fast algorithms for frequent itemset mining using fp-trees. IEEE Trans. Knowl. Data Eng. 17(10), 1347–1362 (2005), http://dx.doi.org/10.1109/TKDE.2005.166
7. Guns, T., Nijssen, S., De Raedt, L.: k-pattern set mining under constraints. Knowledge and Data Engineering, IEEE Transactions on 25(2), 402–418 (2013)
8. Guns, T., Nijssen, S., De Raedt, L.: Itemset mining: A constraint programming perspective. Artificial Intelligence 175(12), 1951–1983 (2011)
9. Han, J., Pei, J., Yin, Y., Mao, R.: Mining frequent patterns without candidate generation: A frequent-pattern tree approach. Data Min. Knowl. Discov. 8(1), 53–87 (2004), http://dx.doi.org/10.1023/B:DAMI.0000005258.31418.83
10. Khiari, M., Boizumault, P., Crémilleux, B.: Constraint programming for mining n-ary patterns. In: CP 2010. pp. 552–567 (2010), http://dx.doi.org/10.1007/978-3-642-15396-9_44
11. Mannila, H., Toivonen, H.: Levelwise search and borders of theories in knowledge discovery. Data Mining and Knowledge Discovery 1(3), 241–258 (1997)
12. Nijssen, S., Guns, T.: Integrating constraint programming and itemset mining. In: Machine Learning and Knowledge Discovery in Databases, European Conference, ECML PKDD 2010, Barcelona, Spain, September 20-24, 2010, Proceedings, Part II. Lecture Notes in Computer Science, vol. 6322, pp. 467–482. Springer (2010)
13. Pasquier, N., Bastide, Y., Taouil, R., Lakhal, L.: Efficient mining of association rules using closed itemset lattices. Inf. Syst. 24(1), 25–46 (1999), http://dx.doi.org/10.1016/S0306-4379(99)00003-4
14. Pei, J., Han, J., Mao, R.: CLOSET: an efficient algorithm for mining frequent closed itemsets. In: SIGMOD Workshop on Data Mining and Knowledge Discovery. pp. 21–30 (2000)
15. Rossi, F., Beek, P.v., Walsh, T.: Handbook of Constraint Programming (Foundations of Artificial Intelligence). Elsevier Science Inc., New York, NY, USA (2006)
16. Uno, T., Asai, T., Uchida, Y., Arimura, H.: An efficient algorithm for enumerating closed patterns in transaction databases. In: DS 2004. pp. 16–31 (2004), http://dx.doi.org/10.1007/978-3-540-30214-8_2
17. Zaki, M.J., Gouda, K.: Fast vertical mining using diffsets. In: SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 326–335 (2003), http://doi.acm.org/10.1145/956750.956788
18. Zaki, M.J., Hsiao, C.: CHARM: an efficient algorithm for closed itemset mining. In: SIAM International Conference on Data Mining. pp. 457–473 (2002), http://dx.doi.org/10.1137/1.9781611972726.27

19. Zaki, M.J., Parthasarathy, S., Ogihara, M., Li, W.: New algorithms for fast discovery of association rules. In: KDD-97). pp. 283–286 (1997), http://www.aaai.org/Library/KDD/1997/kdd97-060.php

Dataset	*minsup* (%)	#Pat	#Propagations		#Nodes		Memory	
			CLOSEDPATTERN	CP4IM	CLOSEDPATTERN	CP4IM	CLOSEDPATTERN	CP4IM
Mushroom	30	428	**1447**	430016	**995**	1039	**18416**	88079
	20	1198	**4250**	1030445	**2881**	3071	**31112**	90831
	10	4898	**17221**	2771719	**11443**	13281	**42840**	93839
	5	12855	**45361**	5574143	**30237**	36495	**69216**	97552
	1	51672	**169781**	13813312	**117851**	168999	**100528**	102608
	0.5	76199	**240618**	18018929	**170091**	259427	111296	**104017**
	0.1	164118	**479630**	31222435	**350491**	529289	143664	**106577**
	0.05	203882	**580333**	36520438	**430799**	622145	148120	**107025**
Chess	60	98393	**314182**	2661395	197063	**196787**	29712	**15623**
	50	369451	**1206361**	9201740	751733	**738907**	31952	**16520**
	40	1366834	**4590519**	30541475	2863847	**2733735**	39720	**18568**
	30	5316468	-	**104618207**	-	**10635019**	-	**19976**
	20	22918586	-	**385399747**	-	**45901933**	-	**19847**
Connect	90	3487	**12582**	1865236	8677	**6973**	**18320**	561872
	80	15108	**59708**	9453279	41939	**30215**	**21752**	585040
	70	35876	**147814**	24968701	105663	**71751**	**22848**	594640
	60	68350	**281560**	51648114	203425	**136699**	**24896**	661583
	50	130102	**538749**	98600221	382441	**260203**	**26968**	665871
Pumsb	95	111	**371**	-	**273**	-	**13600**	-
	90	1467	**4760**	-	**3001**	-	**18416**	-
	85	8514	**30624**	-	**19581**	-	**20920**	-
	80	33296	**137104**	-	**89227**	-	**23280**	-
	75	101048	**432278**	-	**290553**	-	**26688**	-
Retail	10	10	**27**	-	**19**	-	**8760**	-
	5	17	**48**	-	**33**	-	**11096**	-
	1	160	**526**	-	**319**	-	**62568**	-
	0.5	581	**1942**	-	**1161**	-	**425608**	-
	0.1	7696	**25785**	-	**15391**	-	**44152888**	-
	0.05	19699	**65781**	-	**39407**	-	**113664128**	-
T10I4D100K	1	386	**1288**	-	**771**	-	**1207720**	-
	0.5	1074	**3310**	-	**2147**	-	**3118024**	-
	0.1	26807	**72073**	-	**53765**	-	**5074328**	-
	0.05	46994	**132332**	-	**95223**	-	**5519552**	-
	0.01	283398	**935190**	-	**602235**	-	**5947168**	-
T40I10D100K	10	83	**276**	-	**165**	-	**143552**	-
	1	65237	**176994**	-	**130473**	-	**4678544**	-
BMS-Web-View1	0.16	32	**102**	-	**63**	-	**17464**	-
	0.08	9392	**29854**	-	**18935**	-	**1112024**	-
	0.06	64763	**220487**	-	**147811**	-	**1174472**	-
	0.04	155652	**611247**	-	**424203**	-	**1253304**	-
	0.02	422693	**1533714**	-	**1031727**	-	**1489272**	-

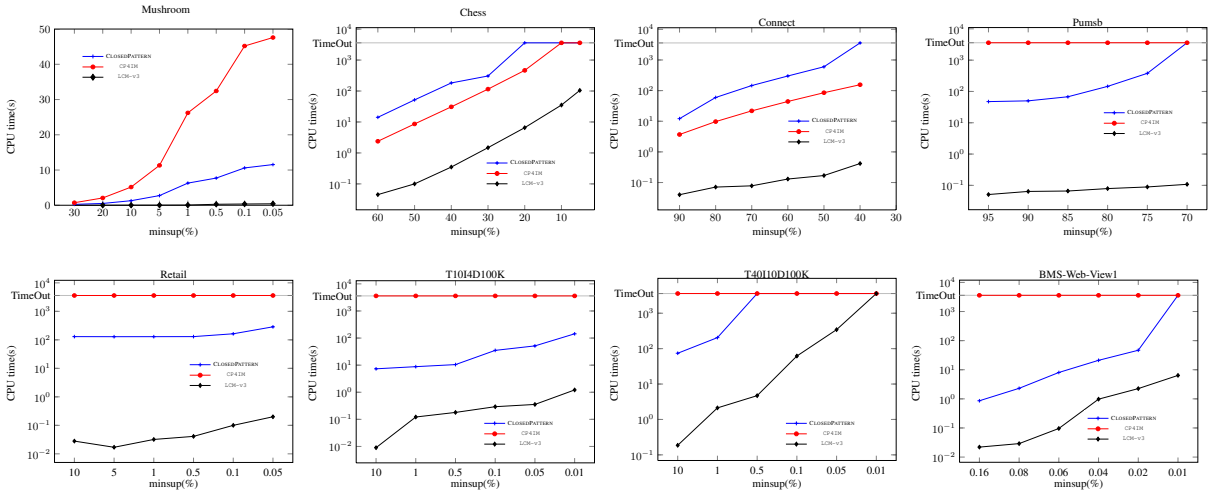

Fig. 1: Computation time on datasets for different values of minimum support.

Table CLOSEDPATTERN vs CP4IM (columns marked by "-" means that the algorithm runs out of memory).